Service Assurance for Virtualized Networks featuring
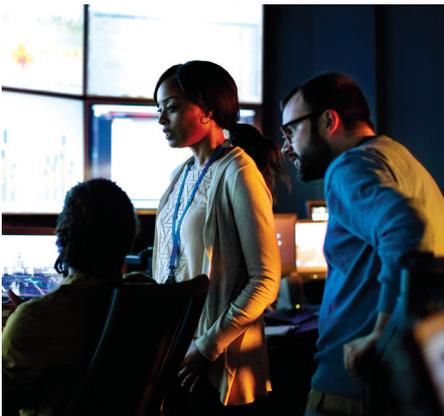Intel® Infrastructure Management Technologies

CPU Performance Counter Monitoring
Intel® Platform Service Assurance

This feature brief describes the telemetry provided by pmu-tools to the collectors on the platform. The collectors provide the statistics to higher-level management systems or telemetry frameworks such as Ceilometer and/or Gnocchi. This information can be used to profile a VNF or application so that it can be tuned for optimal performance.

Feature Description

The CPU Performance Counter Monitoring feature in a collectd environment uses the jevents library (see Ref. 3). The jevents library is a component of the Performance Management Unit (PMU) Tools package, designed specifically for profiling and performance analysis. The information gathered using the jevents library is then made available to higher-level management systems using the intel_pmu plug-in that is part of the telemetry collection framework. The information helps profile applications or VNFs to trace dynamic control flow and identify potential hotspots in support of platform service assurance.

Feature Data Sets

Monitored event categories for the intel_pmu plug-in include:

- HWCacheEvents – Hardware cache events
- KernelPMUEvents – Kernel PMU events
- SWEvents – Software events
- HWSpecificEvents – Hardware-specific events[1]

Configuration

Configuration of the CPU Performance Counter Monitoring feature is done in the intel_pmu plug-in section of the collectd plug-in configuration file. Each event category (cache, PMU, software, and hardware-specific) can be enabled or disabled and the collection interval can be specified.

Telemetry Collection Framework Support

The CPU Performance Counter Monitoring feature uses a collectd plugin called intel_pmu to collect cache, kernel PMU, software and hardware-specific events and provide them to higher-level management systems (see Ref. 4 for details on collectd).

Abstract

The CPU Performance Counter Monitoring feature is a means of gathering key CPU performance metrics for integration with higher-level management systems. Hardware events such as instructions executed, cache-misses suffered, or branch mispredictions are captured in dedicated CPU hardware registers, called performance counters. These performance counters are often used to conduct low-level performance analysis for tuning and provide important information for platform service assurance. The feature uses the Linux pmu-tools (see Ref. 1) to retrieve the information. Pmu-tools is a collection of tools for profile collection and performance analysis on Intel® CPUs on top of the Linux perf utility (see Ref. 2). This utility uses performance counters in the CPU.

[1] To use hardware-specific event names in the configuration file, it is necessary to download the events list file for the current CPU. The events file list for the CPU can be downloaded using the event_download.py script, which is part of the pmu-tools package.

External Interface Support

There is currently no external interface support for this feature.

Feature Dependencies

The CPU Performance Counter Monitoring feature depends on having the following features running on the platform:

- PMU tools (see Ref. 1 for details)
- Jevents (see Ref. 3 for details)
- Collectd must be configured on the platform (see Ref. 4 for details)
- Collectd Ceilometer[2], Gnocchi, and Aodh plug-ins for delivery of metrics to OpenStack* (see Ref. 5 for details)

Where to Get More Information

For more information, visit https://networkbuilders.intel.com/network-technologies/serviceassurance

REFERENCES	
TITLE	LINK
Ref. 1: Pmu-tools	https://github.com/andikleen/pmu-tools
Ref. 2: Perf: Linux profiling with performance counters	https://perf.wiki.kernel.org/index.php/Main_Page
Ref. 3: Jevents	https://github.com/andikleen/pmu-tools/tree/master/jevents
Ref. 4: Collectd	https://collectd.org/documentation.shtml
Ref. 5: OpenStack Ceilometer, Gnocchi, and Aodh	https://wiki.openstack.org/wiki/Telemetry

[2] The Ceilometer interface is being deprecated. Platform telemetry is delivered directly to OpenStack Gnocchi, and events are delivered to OpenStack Aodh. See Ref. 4 for details.